


CONNEXIONS



TM

The Interoperability Report

September 1989

Volume 3, No. 9

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Components of OSI: The Session Service.....	2
White Pages Pilot.....	10
Letter to the Editor.....	15
Effect of the Internet Worm on Security.....	16
Upcoming Events.....	19
AppleTalk versus IP.....	20
IAB Reorganized.....	25

ConneXions is published monthly by
Advanced Computing Environments,
480 San Antonio Road, Suite 100,
Mountain View, California 94040, USA.
Phone: 415-941-3399. Fax: 415-949-1779.

© 1989
Advanced Computing Environments.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report
and the *ConneXions* masthead are
trademarks of Advanced Computing
Environments.

ISSN 0894-5926

From the Editor

As we continue our series *Components of OSI*, with an article on the OSI Session Service, I cannot resist telling you one of the very few OSI jokes I know. It is attributed to the *Cynic of the Internet*, Dr. Paul V. Mockapetris of USC-ISI, and goes like this:

Question: "What is the difference between Session and a bus?"
Answer: "Sometimes you miss the bus..."

Hopefully, after reading the article by Kim Banker of Hewlett Packard, you will discover that the above is just a joke—and that the OSI Session Service indeed provides a wide range of capabilities which applications can build upon. Contrast this to TCP/IP, where each application duplicates the session functionality, rather than relying on a common service.

It is encouraging to see early implementation and experimentation with the emerging OSI protocols take place in existing networking environments such as the Internet. The NYSERNet *White Pages Pilot* represents a large-scale experimental deployment of the CCITT/ISO X.500 Directory system. In an article on page 10 we bring you an overview of this project. The overview is based on an article which appeared in the July/August 1989 issue of *NYSERNET News*. (Vol. 2, No. 8).

Security continues to be an important topic for the networking community, particularly in the wake of the infamous "Internet Worm" of last November. Fred Ostapik of SRI International discusses some security issues in an article on page 16. SRI and Advanced Computing Environments will be hosting a seminar on computer and network security, called *InfoSec*, at the end of November, 1989. Information on InfoSec can be found on page 19.

Greg Minshall compares the AppleTalk "plug-and-play" networks to IP networks which require at least one "Wizard" in order to install and operate. Greg argues that much could be done to improve this state of affairs.

The Internet Activities Board (IAB) was recently reorganized into two main task forces; the *Internet Engineering Task Force* (IETF) and the *Internet Research Task Force* (IRTF). The "new IAB" is described on page 25.

INTEROP™ 89 is only one month away! Be sure to call in your registration now at 415-941-3399. See you in San Jose!

Components of OSI: The Session Service

by Kim Banker, Hewlett Packard Corporation

Introduction

This article provides an overview of the OSI Session service as specified in the joint ISO/CCITT international standards ISO IS 8326, CCITT X.215 [1] and ISO IS 8327, CCITT X.250 [2]. The Session service definition and protocol specification became an officially approved international standard in 1984. After five years, it is still perhaps one of the most misunderstood and misinterpreted ISO standards. Relatively few people fully understand the capabilities and limitations the Session service provides. This article outlines the basic services of the Session layer and points out many of the misunderstandings that plague this ISO protocol.

Figure 1 illustrates how the Session layer fits into the OSI environment. Through services provided by the Transport layer, the *Session service provider* (SS-provider) furnishes the binding relationship and dialogue services between two cooperating *Session service users* (SS-users). The SS-provider is that entity within a protocol stack implementation that performs the Session services for the SS-user. Although the presentation service provider is the most probable SS-user to utilize the Session services, some applications (for example, X.400, version 1984) access Session services directly.

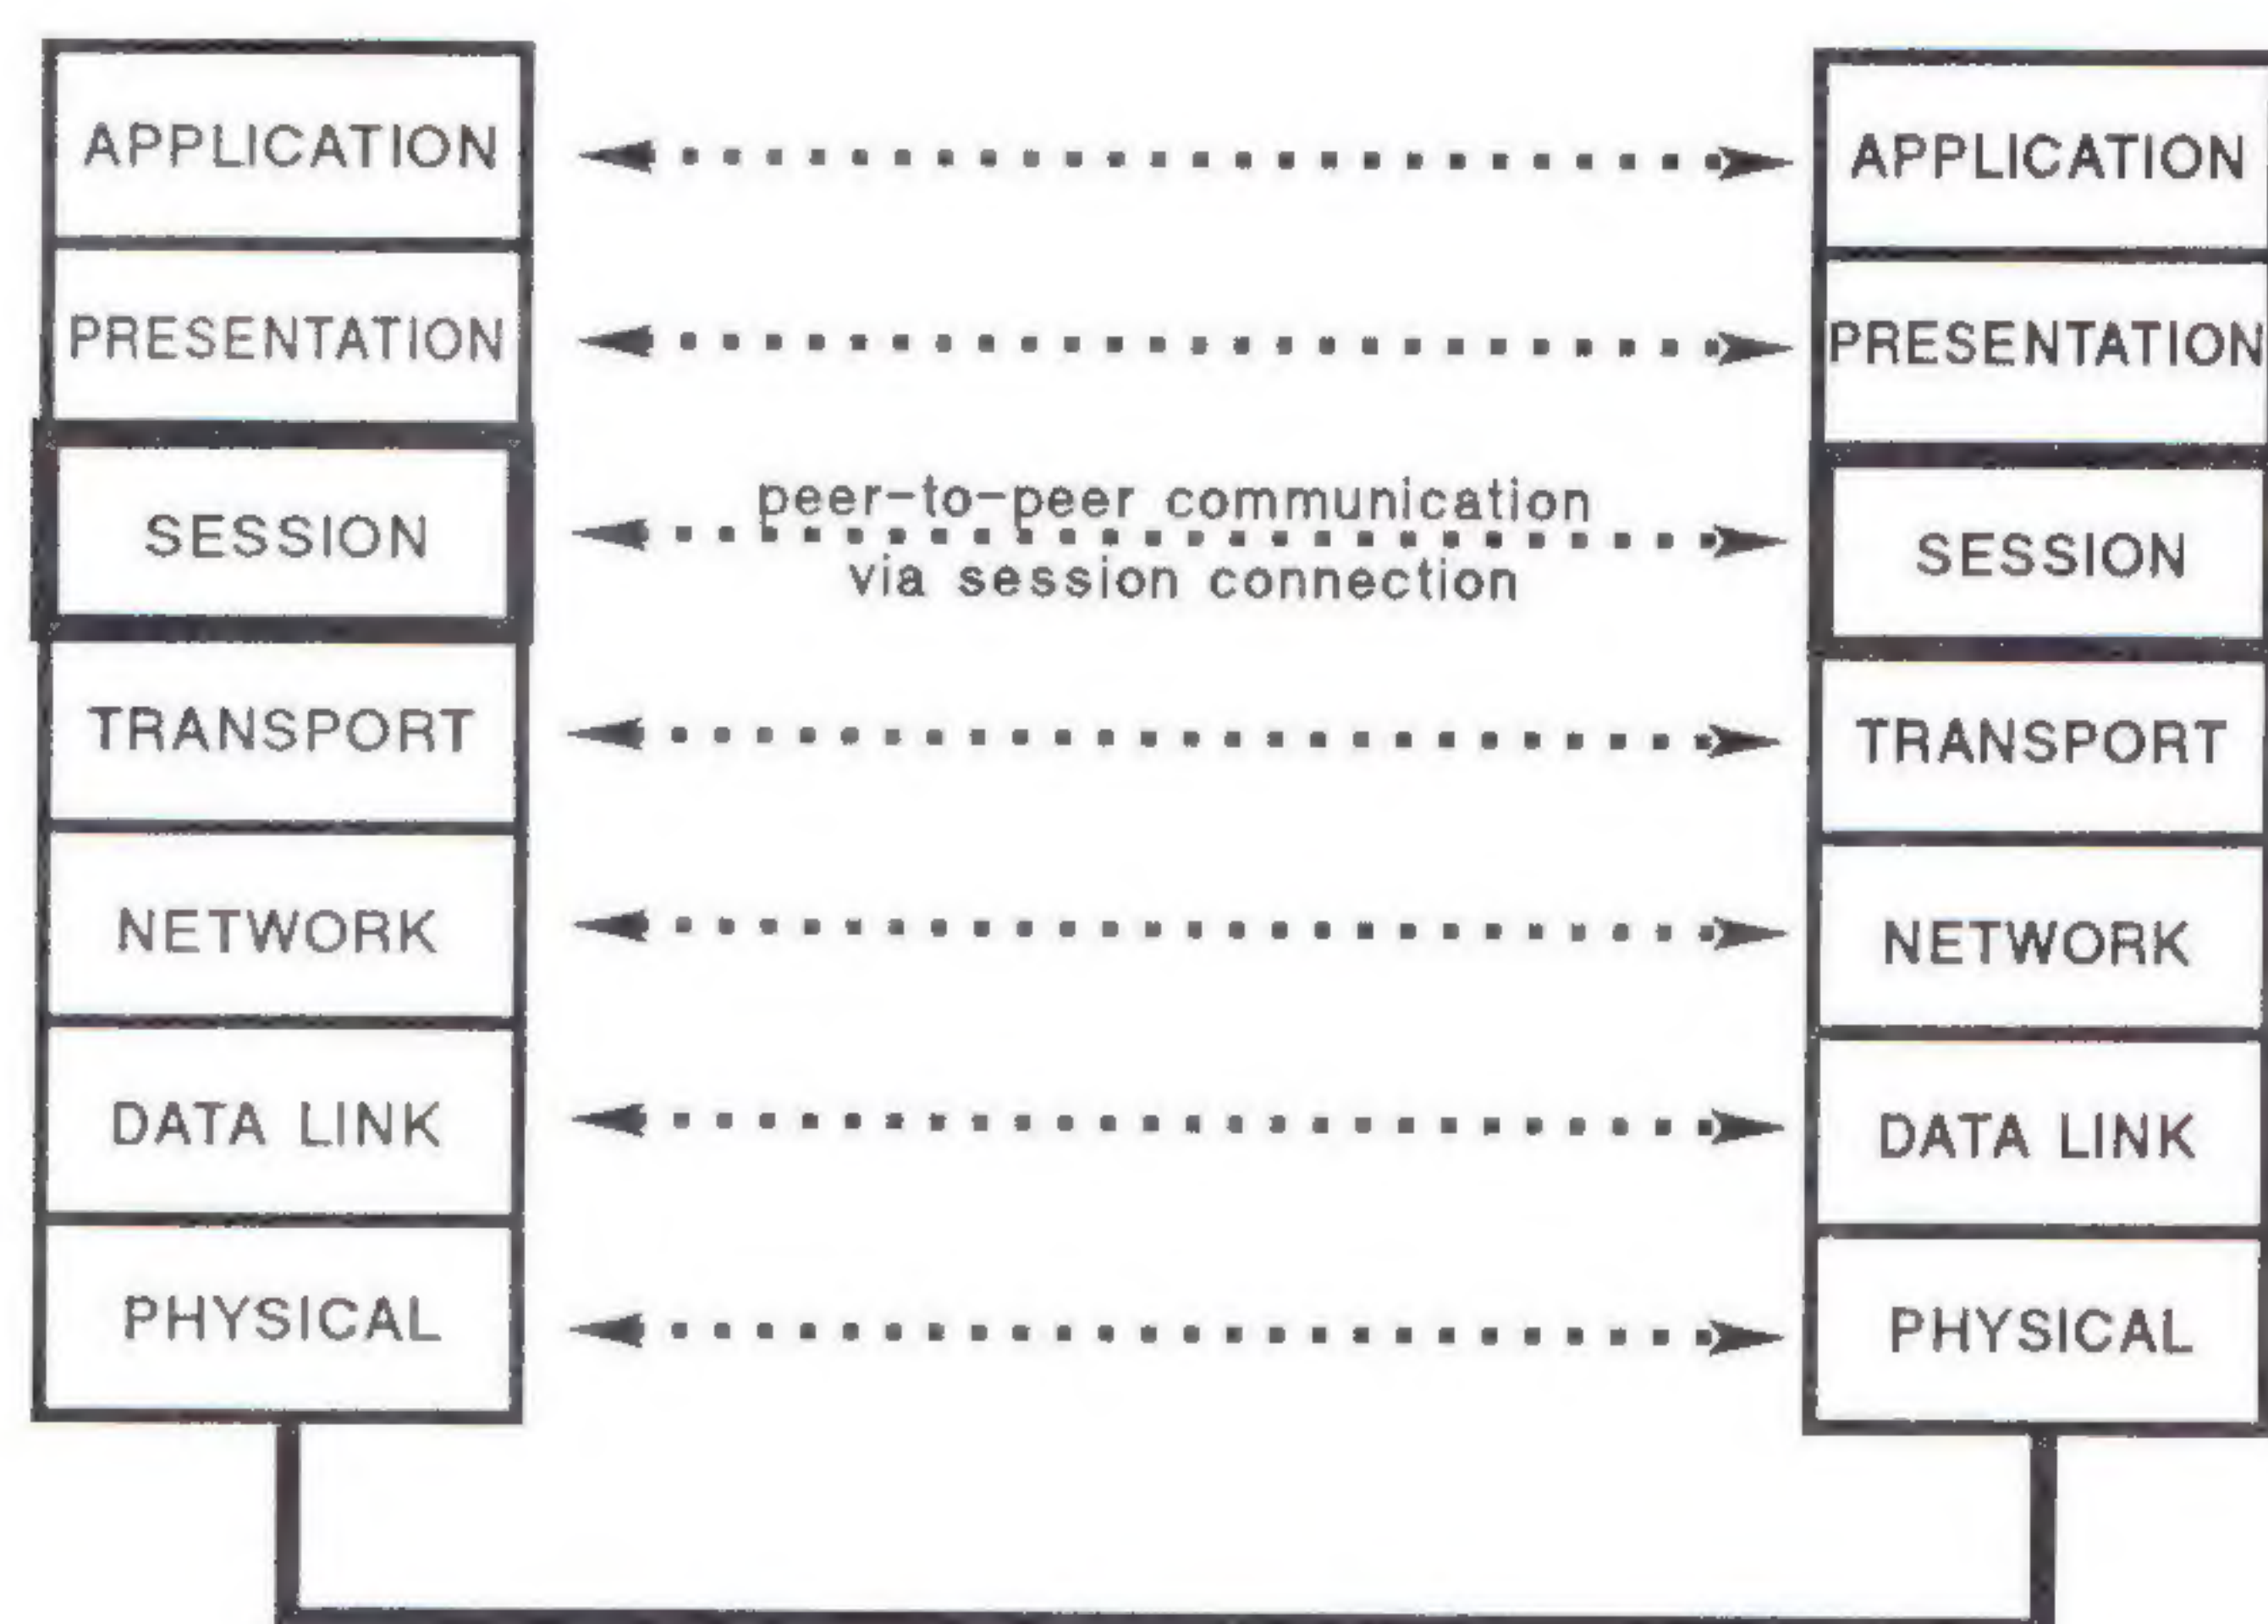


Figure 1: Open Systems Interconnection (OSI) Environment

The binding created between two SS-users is called a *Session Connection*. Despite a recently adopted ISO Session addendum [5] to provide for a "connectionless" Session service, the base Session standard defines Session as a connection oriented protocol. Like most connection oriented protocols, a Session connection is analogous to making a telephone call. Once a communication channel is established between two users, information and services are exchanged, and then the communication channel is disconnected. It is these same three phases that comprise a Session connection. Let's examine each of the Session connection phases in detail and the associated services that are provided to SS-user.

Session Connection Establishment Phase

The connection establishment phase permits two peer SS-users to establish the communication link that will give access to a number of useful Session services. Simultaneous Session connections may be opened, only limited by the internal resources of the SS-provider.

The calling SS-user (requester of the Session connection) issues a connection request to the called SS-user (responder of the Session connection). The called SS-user then has the option of either accepting or refusing the Session connection and sends the appropriate protocol response. During this protocol exchange, a number of parameters are exchanged between the two SS-users, and a series of negotiations occur to establish the characteristics of the Session connection. Just as in making a telephone call, addressing and directory information must be exchanged to establish the locations of the two SS-users. For the Session layer, SSAPs (*Session Service Access Points*) define the source and destinations of the two SS-users. Once the Session connection has been established, SSAP information is no longer required to be exchanged.

Among the negotiations that occur during connection establishment, functional units and initial token settings are among the most significant. *Functional units* define the various capabilities available to the SS-users during any one given Session connection. Ten different functional units are available in a fully functional Session implementation. The services associated with the functional units will be described in more detail when we discuss the data transfer phase of the Session connection. Each SS-user indicates which of the ten functional units they desire for the Session connection, and only the common intersection of the two selections become the operating functional units that will be used during that Session connection.

Tokens

Session uses “tokens” to grant control of certain Session capabilities to one of the two SS-users at a time. The initial possession of the tokens is established during connection establishment and then during the lifetime of the Session connection the tokens can be exchanged according to the requirements of the SS-users. Four different tokens may be available to the SS-users, depending upon the functional units that have been negotiated:

- data token
- synchronize-minor token
- major/activity token
- release token

The use of these tokens is described in more detail as we discuss the Session services that use the individual tokens.

In addition to other negotiations, such as version number, initial synchronization number, and segmentation, each SS-user is permitted to send “user data” on the Session connect request and response. User data is entirely transparent and meaningless to the SS-provider, who merely passes the unaltered information from one SS-user to the other. When the original Session standard was designed the majority of user information was expected to be sent on Session’s data transfer services, and the amount of user data on Session’s other service primitives was limited to 512 octets of data. However, since then, many application protocols have been designed to “piggyback” their protocol data units inside the user data fields of Session protocol data units, in particular the connect and accept SPDUs. With the recent approval of an unlimited data addendum [4] to the Session service and protocol, user data is no longer limited in size and is available on all Session service primitives.

continued on next page

Components of OSI: The Session Service (*continued*)

Data Transfer Phase

Once the connection is established the more significant communication tasks can commence and the Session connection enters what is termed the "Data Transfer State." As the name implies the bulk of data is expected to be exchanged during this phase of the connection, but much more than data transfer takes place while the Session connection is open.

Token Exchange

As mentioned during the description of the connection establishment phase, once the initial possession of the Session tokens has been established, any or all of the four tokens that are available can be exchanged between the two SS-users. An SS-user may ask for any available token that it does not own with a "please token" service primitive. Likewise, an SS-user may give away any token that it owns to the other SS-user with a "give token" service primitive. This simple exchange of token control provides much of the locking and control mechanisms required by certain applications.

The most common service the Session layer provides is the transfer of data from one SS-user to the other. Session provides this service in four different flavors...actually five counting "capability data," which will be discussed under the topic of activities.

Normal Data Transfer

Normal data transfer can operate in one of two modes—full or half duplex, but not both. The decision to operate in either of these modes is negotiated during connection establishment and cannot be altered during the lifetime of the connection. Full duplex (also known as Two-Way-Simultaneous) permits both SS-users to send data freely in both directions at the same time, independent of any token restrictions. Half duplex (Two-Way-Alternate) makes use of the data token by restricting the data transmission capability to one SS-user at a time. Only the owner of the data token may transmit data. The data token must then be exchanged, using the token exchange services, to permit the other SS-user to send data.

Expedited Data Transfer

Expedited data permits an SS-user to transmit a very small (14 octets maximum) amount of user data that is sent free from the token and flow control constraints associated with the normal data transfer service. The expedited data functional unit must be negotiated by both SS-users to grant this service to the SS-user.

Typed Data Transfer

Typed data allows SS-users to send information independent of the availability and assignment of the data token. This might seem to contradict the intention of half-duplex data transfer, but when used properly between two cooperating SS-users typed data can provide for a powerful mixed half/full duplex mode that many applications today are finding useful. Since typed data can only be sent with the typed data functional unit, it is still possible to restrict data transfer to purely half duplex by not negotiating the typed data functional unit during connection establishment.

Synchronization

One of the original goals in defining the Session layer was to provide the functions of "dialogue control" for the application layer protocols. One of the characteristics of a dialogue is to break up an exchange of data into a series of distinct and separate units called *dialogue units*. The *synchronization* services, together with the *resynchronization* and *activity* services, provide much of the Session dialogue control capabilities.

Associated with all the synchronization and resynchronization services are *serial numbers*. A valid serial number, ranging from 0 to 999999, defines synchronization points that are inserted between the data exchanges of SS-users. These inserted synchronization points can then be used as reference points along the data stream for future access or recovery purposes. The SS-provider maintains and regulates the serial numbers for the SS-users, however, any semantics which the SS-users give to the synchronization points is entirely transparent and meaningless to the SS-provider. From the original negotiation of the serial number at connection establishment, the SS-provider increments the serial by one for each synchronization point established by the SS-users.




Figure 2: Structured Dialogue Unit

The original design of Session synchronization is two-way-alternate in nature. That is, a token is used to control which SS-user is permitted to insert synchronization points into the data flow. A recently approved symmetric synchronization addendum [3] to the Session service and protocol allows for two-way-simultaneous synchronization without the need for token restrictions. This contribution to the Session service permits SS-users to perceive of two independent directions of data flow instead of one. There are two classifications of synchronization: Major synchronization and Minor synchronization

Major synchronization

Major synchronization points determine the starting and ending points of dialogue units. As shown in figure 2, when an SS-user determines the need to start a new dialogue unit (at the same time ending any current dialogue unit) a major synchronize request made to the SS-provider will indicate the next serial number to both SS-users simultaneously. Establishing a major synchronization points along a data flow implies that all data sent/received prior to that point has been properly processed and accounted for by both SS-users, never requiring a recovery or retransmission of that data.

Minor synchronization

Minor synchronization points are used to structure the exchange of data within a dialogue unit, again according to semantics defined by the SS-users. Serial numbers are indicated to both SS-users in the same manner as in major synchronization, but as the name implies the consequences of minor synchronization are not as significant as major synchronization. Although the difference between major and minor synchronization seems trivial, the effects on resynchronization are significant, as we'll soon discover.

continued on next page

Components of OSI: The Session Service (continued)

It's important to understand that for both major and minor synchronization the SS-provider only maintains and manages the distribution of the serial numbers to the SS-user. The SS-provider makes no association of the serial numbers to any data exchanged among the SS-users. It is the SS-users responsibility to properly incorporate the serial numbers within the context of their application.

Resynchronization

One of the real value-added services that Session provides is *resynchronization*. The basic concept behind resynchronization is to assist in orderly reestablishment of communication and setting the Session connection to a defined stable state, so that both SS-users are "in synch" and operating from the same point of reference.

Once synchronization points have been established in a data flow three resynchronization options are provided to the SS-user, as illustrated in Figure 3.

Abandon

Abandon is used to resynchronize the Session connection to a point beyond than the current synchronization point. This, in effect, implies to both SS-users that all data sent prior to the abandon resynchronization will be lost or discarded by each SS-user. A resynchronize abandon following the minor sync at number 2036 in Figure 3, for example, would reestablish the current serial number to any point, defined by the SS-user, greater than serial number 2036. The SS-users would then lose future capability to resynchronize to any synchronization points prior to and including serial number 2036.

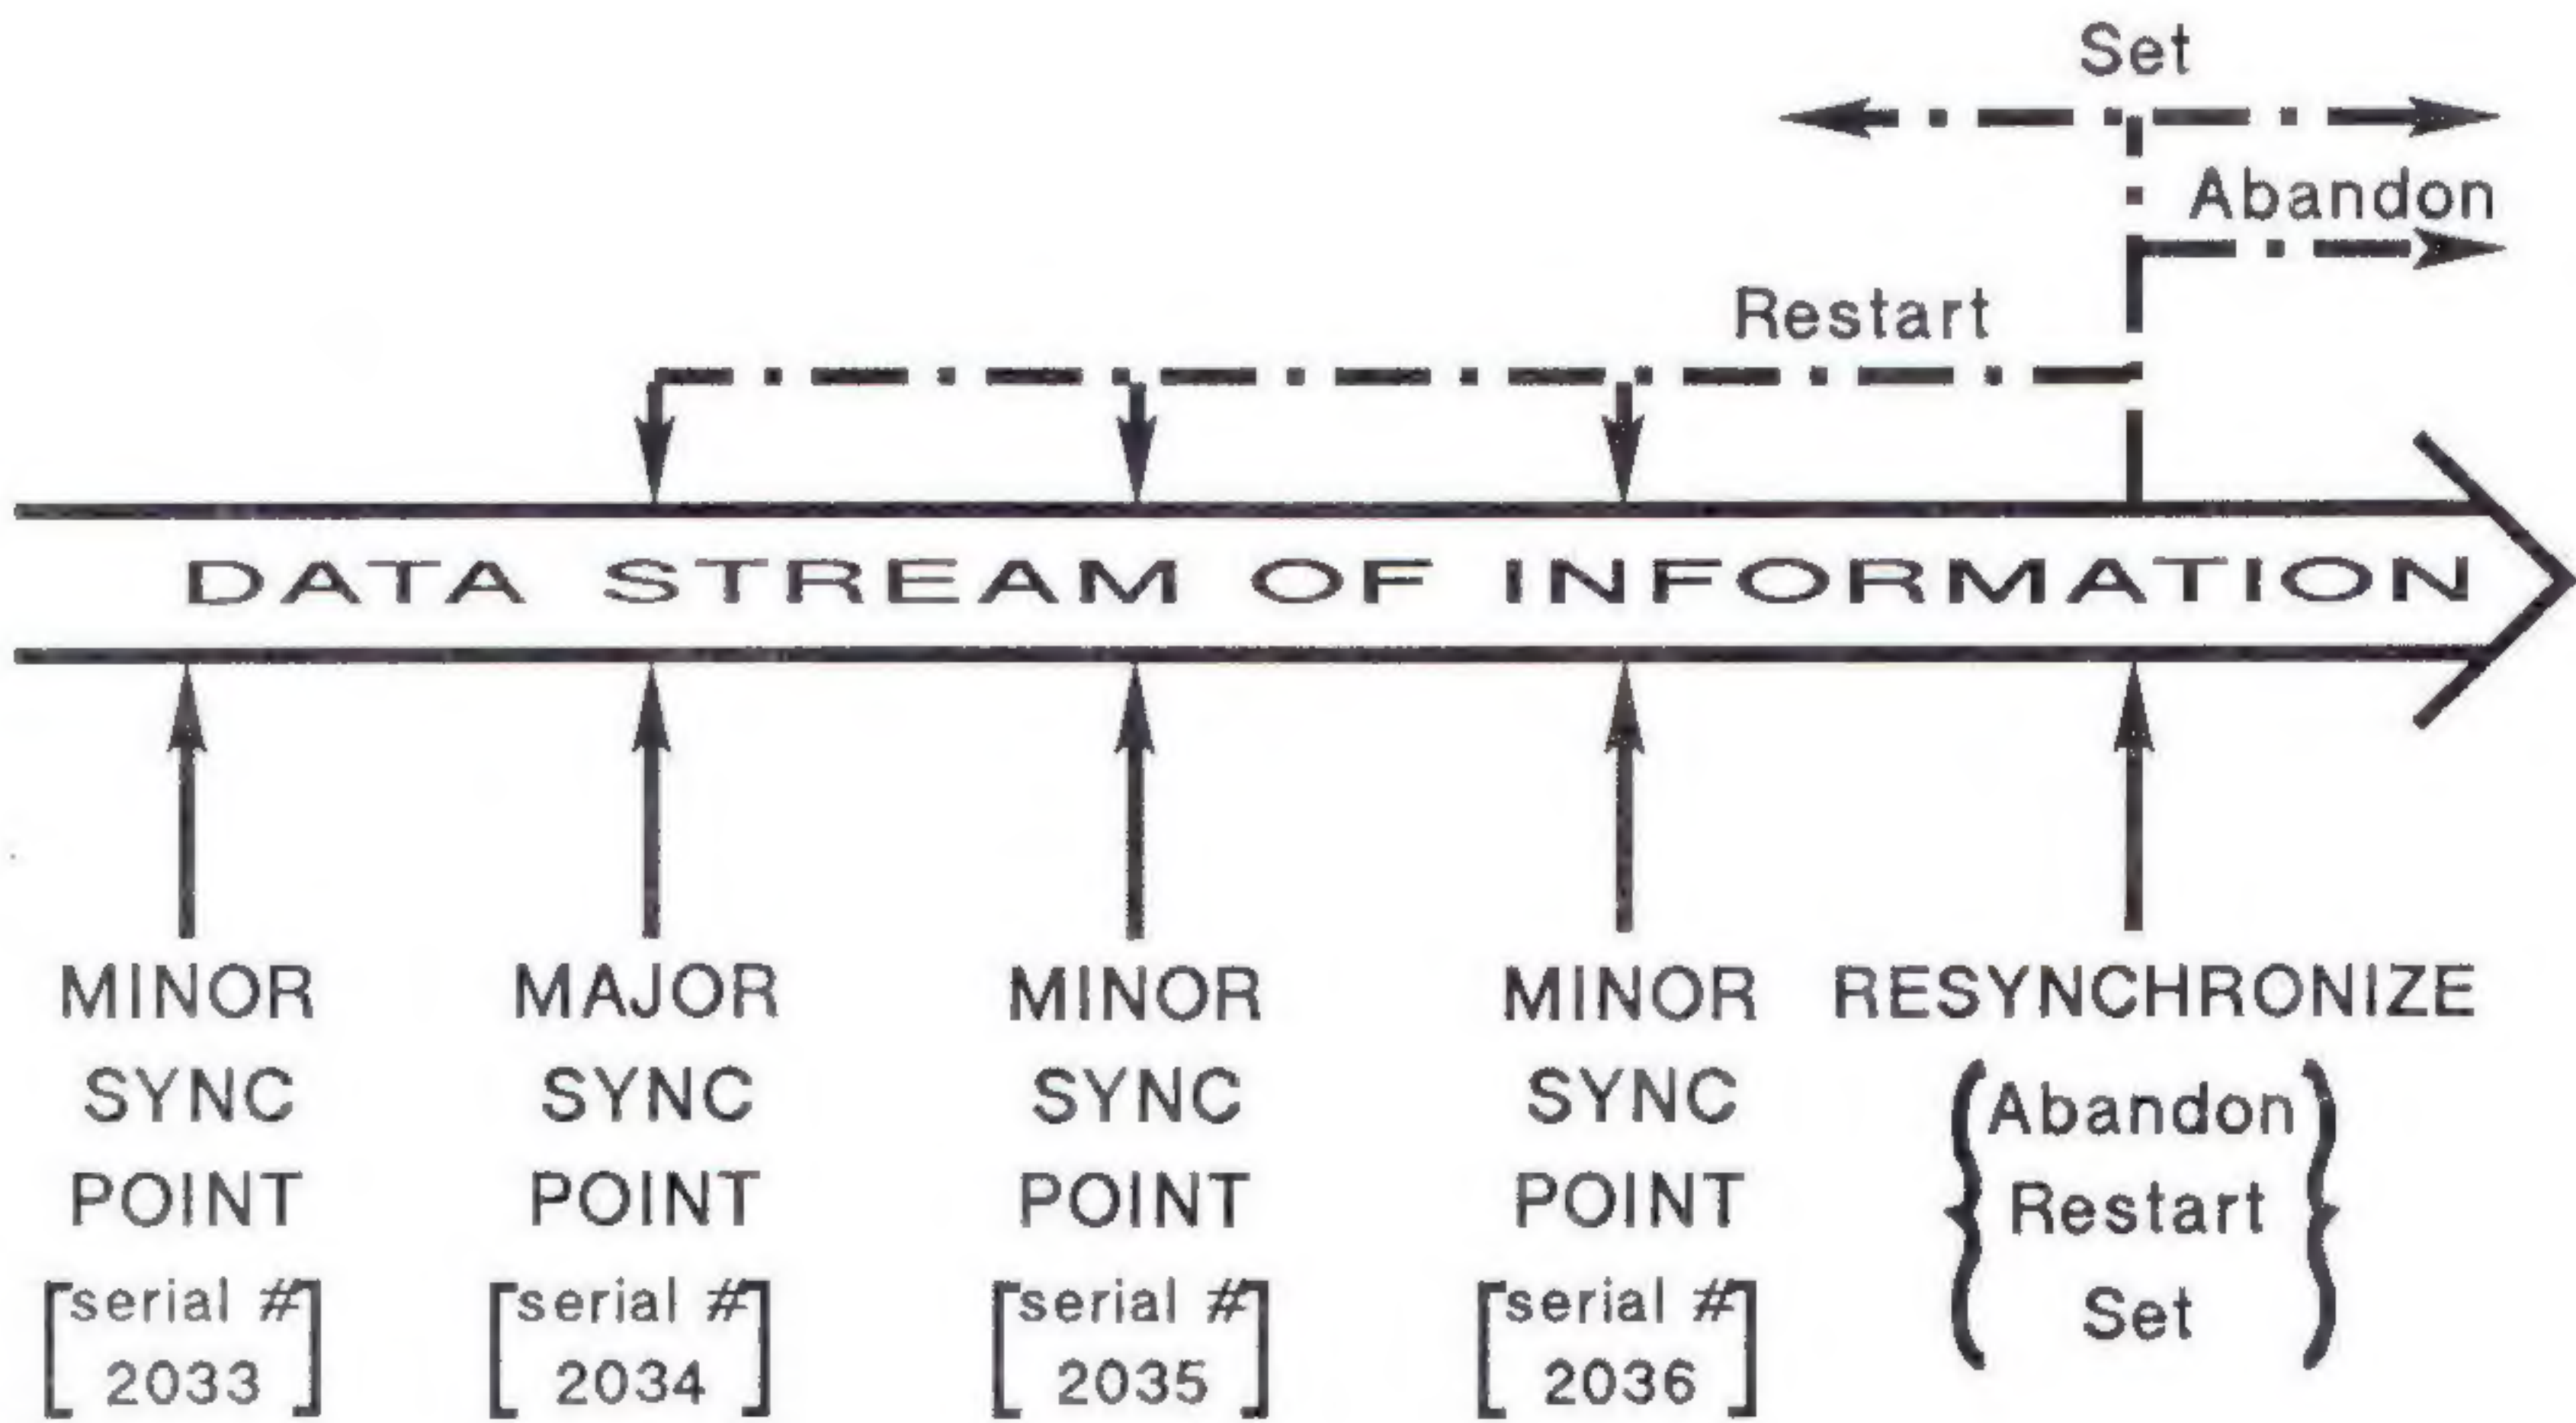


Figure 3: Resynchronization Options

Restart

Restart is used to return to an agreed point which is identified by a previously established synchronization point serial number. This point cannot be any earlier than the last major synchronization point (signifying the primary distinction between major and minor synchronization). In our Figure 3 example, a restart synchronization would permit the SS-users to "return" to synchronization point numbers 2036, 2035, or 2034, but not to 2033. The important thing to note here is that it is the SS-users responsibility to restore all data and semantics associated with the synchronization points following a resynchronization.

A common misunderstanding is that the SS-provider will magically be able to restore any previously sent data during a resynchronization, when in fact the SS-provider only manages the mechanics of the serial numbers and assures both SS-users remain in synch.

Set The final resynchronization option, *set*, is used to synchronize to any valid synchronization point specified by the SS-users. One can think of the set option as a more general case of the abandon option, allowing the SS-users unrestricted choice of the next synchronization point serial number. As in abandon, all rights to resynchronize to points prior to a set resynchronization are lost.

Activities The concept of an *activity* is to group together one or more dialogue units into separate distinguishable pieces of work. As with dialogue units, the semantics associated with activities is entirely the responsibility of the SS-users. Only one activity is allowed on a Session connection at time, but there may be several consecutive activities during any given Session connection. An activity may also span more than one Session connection. An activity can be interrupted and then resumed on the same or on a subsequent Session connection.

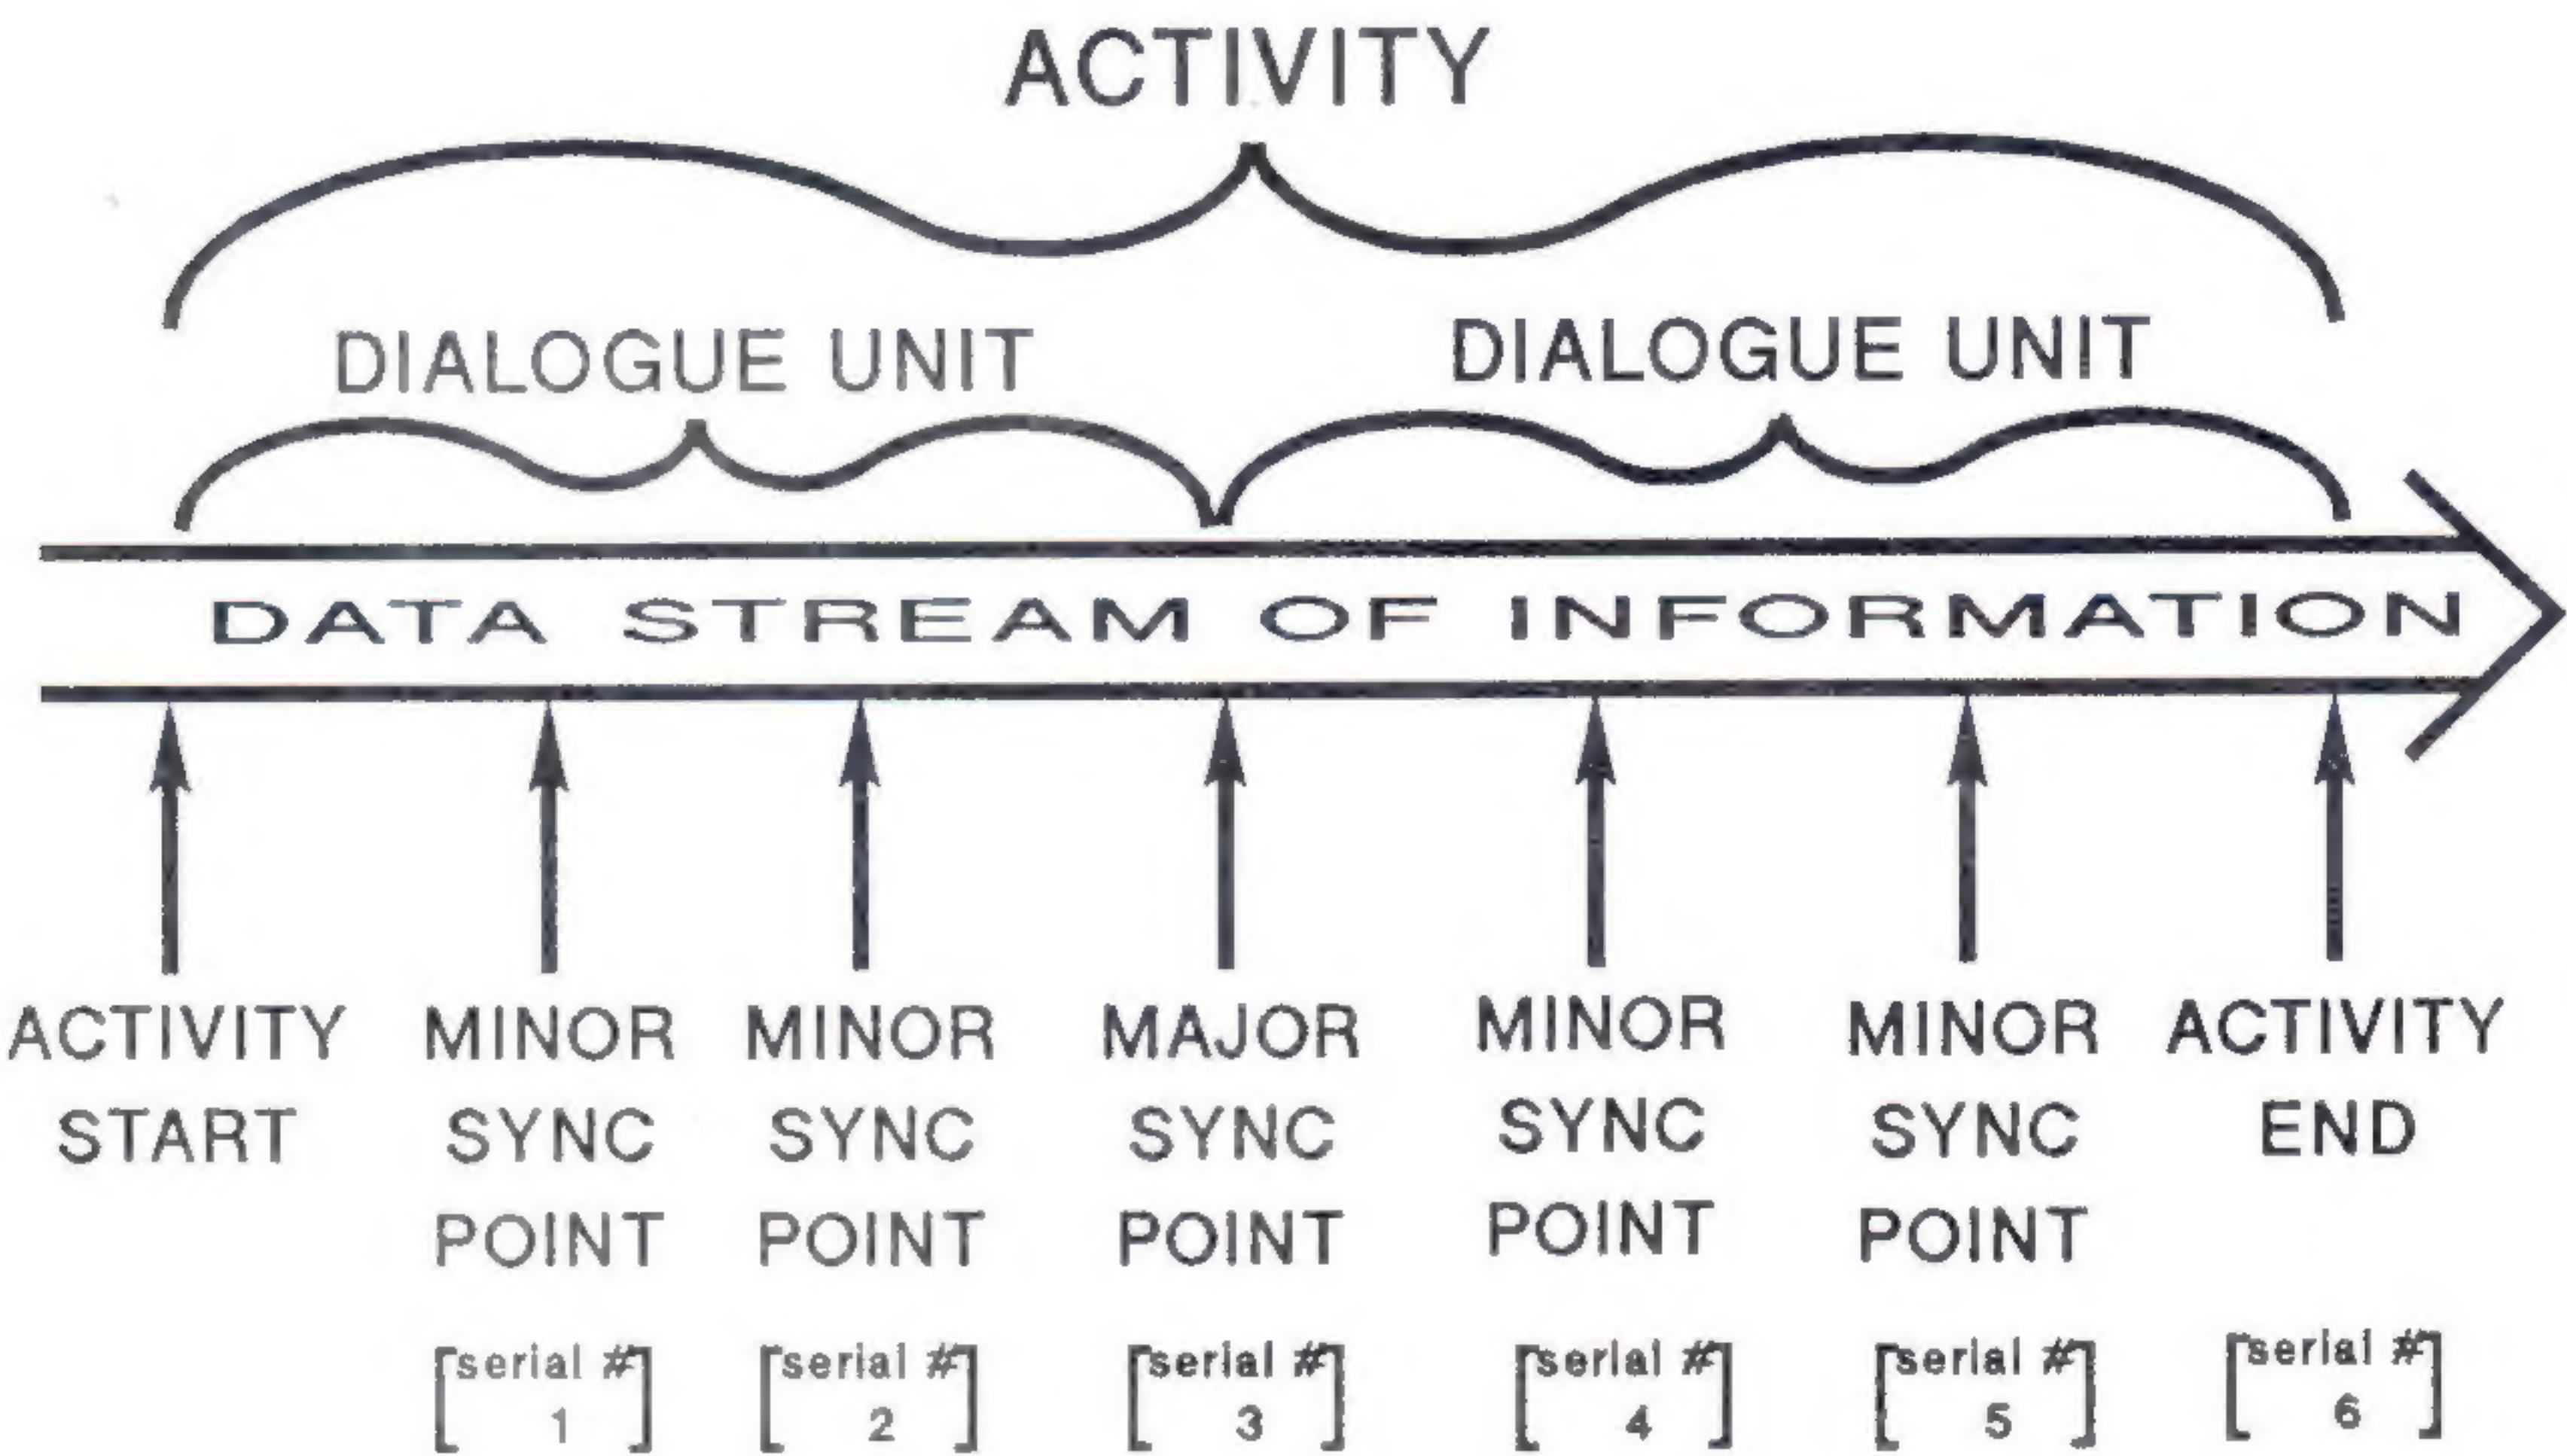


Figure 4: Structured Activity

Figure 4 shows the structure of a Session activity consisting of dialogue units that are separated through the use of major synchronization points. The activity management services provided by the SS-provider include:

Activity Start The *activity start* service initiates a new activity on an already established Session connection. The value of the next synchronization point serial number is automatically set to the value of 1 and cannot be negotiated among the SS-users. An activity can only be started if no activity is in progress and only by the SS-user who has control of the Session tokens.

Activity End To gracefully terminate an activity, the *activity end* service primitive is provided. This service also has the effect of setting a major synchronization point. Like the activity start, only the SS-user who has control of the Session tokens may request an activity end. Terminating an activity does not close the underlying Session connection.

Components of OSI: The Session Service (*continued*)

Activity Discard	In the event an activity needs to be abnormally terminated the <i>activity discard</i> service is available. The implied meaning to the SS-user is that the previous content of the activity is lost and all work associated with the activity will never continue.
Activity Interrupt	At the discretion of the SS-users, an activity may be temporarily terminated, with the intention of continuing the activity at a later time. This <i>activity interrupt</i> implies that the work completed up to that time is not to be discarded and may be resumed later.
Activity Resume	The <i>activity resume</i> service allows as SS-user to indicate that a previously interrupted activity (by the activity resume service) is resumed. It is the SS-users responsibility to restore all information and semantics associated with the activity, this includes the new synchronization point serial number that is to be used for the resumed activity. The resumed activity need not reside on the same connection with which it was interrupted.
Capability Data Transfer	Although not actually part of the activity management function unit, <i>capability data</i> is an integral part of the activity service, and in fact can only be used when activity management has been negotiated for the Session connection. Capability data is used to exchange data when the Session connection is not currently within an activity. Unlike the other data transfer services (normal, typed, and expedited), capability data is a confirmed service with an acknowledgement sent from the receiver of the data back to the sender.
<p>As you might sense from the brief descriptions above, Session activities create a fairly restrictive environment for the SS-user. For example, synchronization and resynchronization is only permitted outside of an activity. These and similar restrictions have created some controversy over the use of activities in certain applications. The activity services were primarily intended for message handling applications like CCITT X.400, and for those services the activity environment is well suited. Some experts, however, feel that the activity service definition is outside the scope of Session services since it behaves in many ways as a sub-layer above the Session layer.</p>	
Exceptions	Two Session services are provided for the reporting of errors or unanticipated situation. These types of errors can be considered "soft errors" with the possibility of recovery.
Provider Exception Reporting service	The <i>provider-initiated exception reporting</i> service permits SS-users to be notified of abnormal conditions within the SS-provider, such as protocol errors. Although the Session connection is not terminated, it is placed into a state of awaiting recovery where only specific SS-user or SS-provider actions will return the Session connection to a normal data transfer state.
User Exception Reporting service	Similar to the provider exception report, the <i>user-initiated exception reporting</i> service is used and initiated by the SS-user to report exception conditions. The effects on the Session connection are the same, placing the connection into a state of awaiting recovery.
Connection Termination Phase	A Session connection can terminate in one of two manners, either through an orderly release or through one of two Session abort services.

The *Session release* service allows the two SS-users to terminate the Session connection in a cooperative manner without the loss of user data. If the release token was negotiated during connection establishment only the owner of this token is permitted to initiate the release, otherwise either side may do so.

Aborts may be initiated either by the SS-user (*U-abort service*) or by the SS-provider (*P-abort service*). User abort service provides the means by which either SS-user may abruptly terminate the Session connection. Use of this service may result in lost data that was in transit at the time of the abort. Provider aborts are generated by the SS-provider in the event of protocol errors or other internal reasons. Both SS-users are notified of the abort, and like user aborts undelivered data may be lost.

Conclusion The Session service provides a wide range of capabilities upon which applications can be built. It's surprising, however, how much these services are underutilized—Session synchronization and resynchronization are good examples of this. Many application layer protocols today are duplicating much of the functionality that Session provides. Hopefully in the years to come, as application designers become more familiar with ISO Session, the full potential of the Session services can be realized.

- References**
- [1] Information Processing Systems—Open Systems Interconnection “Basic Connection Oriented Session service Definition,” International Standard ISO 8326, August 1987.
 - [2] Information Processing Systems—Open Systems Interconnection “Basic Connection Oriented Session Protocol Specification,” International Standard ISO 8327, August 1987.
 - [3] Information Processing Systems—Open Systems Interconnection “Basic Connection Oriented Session service Definition—Addendum 1: Session Symmetric Synchronization for the Session service,” International Standard ISO 8326/ADD 1, November 1988.
 - [4] Information Processing Systems—Open Systems Interconnection “Basic Connection Oriented Session service Definition—Addendum 2 to Incorporate Unlimited User Data,” International Standard ISO 8326/ADD 2, June 1988.
 - [5] Information Processing Systems—Open Systems Interconnection “Basic Connection Oriented Session service Definition—Addendum 3: Connectionless-Mode Session service,” International Standard ISO 8326/ADD 3, March 1988.

KIM BANKER is a software development engineer at Hewlett Packard Roseville Networking Division. He received his M.S. degree in Computer Engineering at Carnegie-Mellon University in 1979. After working briefly at Bell Laboratories on a large data networking system, Kim joined Hewlett Packard to work in the area of manufacturing automation and computer video technology. For the past 5 years he has been extensively involved in the area of OSI protocols, in particular the Session layer protocol. He currently serves as the US representative and editor on ISO Session and is an ISO Session Rapporteur for ISO/IEC JTC 1/SC21/WG6. He is currently part of a project team at Hewlett Packard developing ISO networking products for business and manufacturing environments.

NYSERNet Sponsors White Pages Pilot

In July of 1989, NYSERNet began offering a white pages service to its members. This service is believed to be the largest pilot project of its kind to offer white pages service using the *OSI Directory Service* (X.500) [1]. Further, the pilot project also represents the first large scale use of OSI services in the Internet.

What are White Pages?

Computer networks form the *infrastructure* between the users they interconnect. For example, the electronic mail service offered by computer networks provides a means for users to collaborate towards some common goal. In the simplest cases, this collaboration may be solely for the dissemination of information. In other cases, two users may work on a joint research project, using electronic mail as their primary means of communication.

Most network services are based on the implicit assumption that each user can supply *infrastructural information* to facilitate information transfers through the network. For example, electronic mail services expect that an originator can supply addressing information for all the intended recipients. It is not necessarily the task of electronic mail, per se, to provide this infrastructural information to the user.

This model works fine in small environments, particularly those where infrastructural information is not difficult to obtain and remember. However, the model does not scale well. Consider the case when the membership of a network consists of hundreds of thousands of users belonging to thousands of organizations. It is no longer reasonable for a single user to provide this information, except in very limited circumstances. Further, it is likely that some of the information changes frequently, due to personnel and other resource movement. The goal of a *white pages* service is to provide the necessary information, and to mask the complexity of the infrastructural information.

Existing Facility

The Internet community is currently served by the WHOIS facility. This is a simple, text-based service originally deployed in 1982. Although the users are predominantly humans, information on some networks, hosts, and so on, are also kept in the WHOIS facility. Currently, it is estimated that there are over 70,000 WHOIS entries.

Although the existing facility has proven useful for many years, many believe that the explosive growth of the Internet has made the existing mechanism unworkable for three reasons: first, because the facility is centralized, there are concerns over the availability of the service. Further, centralization requires that all updates must be inserted by a clerk, which increases the likelihood of the information being out of date.

Second, the facility contains only limited, unstructured information, such as postal and electronic mail addresses. While these are useful, the needs of the community have grown much larger. For example, it is often useful to address correspondence to an organizational role (e.g., "Chair of the Department"). While it is possible for the textual information annotated to each entry to contain such information, given the current informational framework, it is not possible to search for or otherwise mechanically process this attribute.

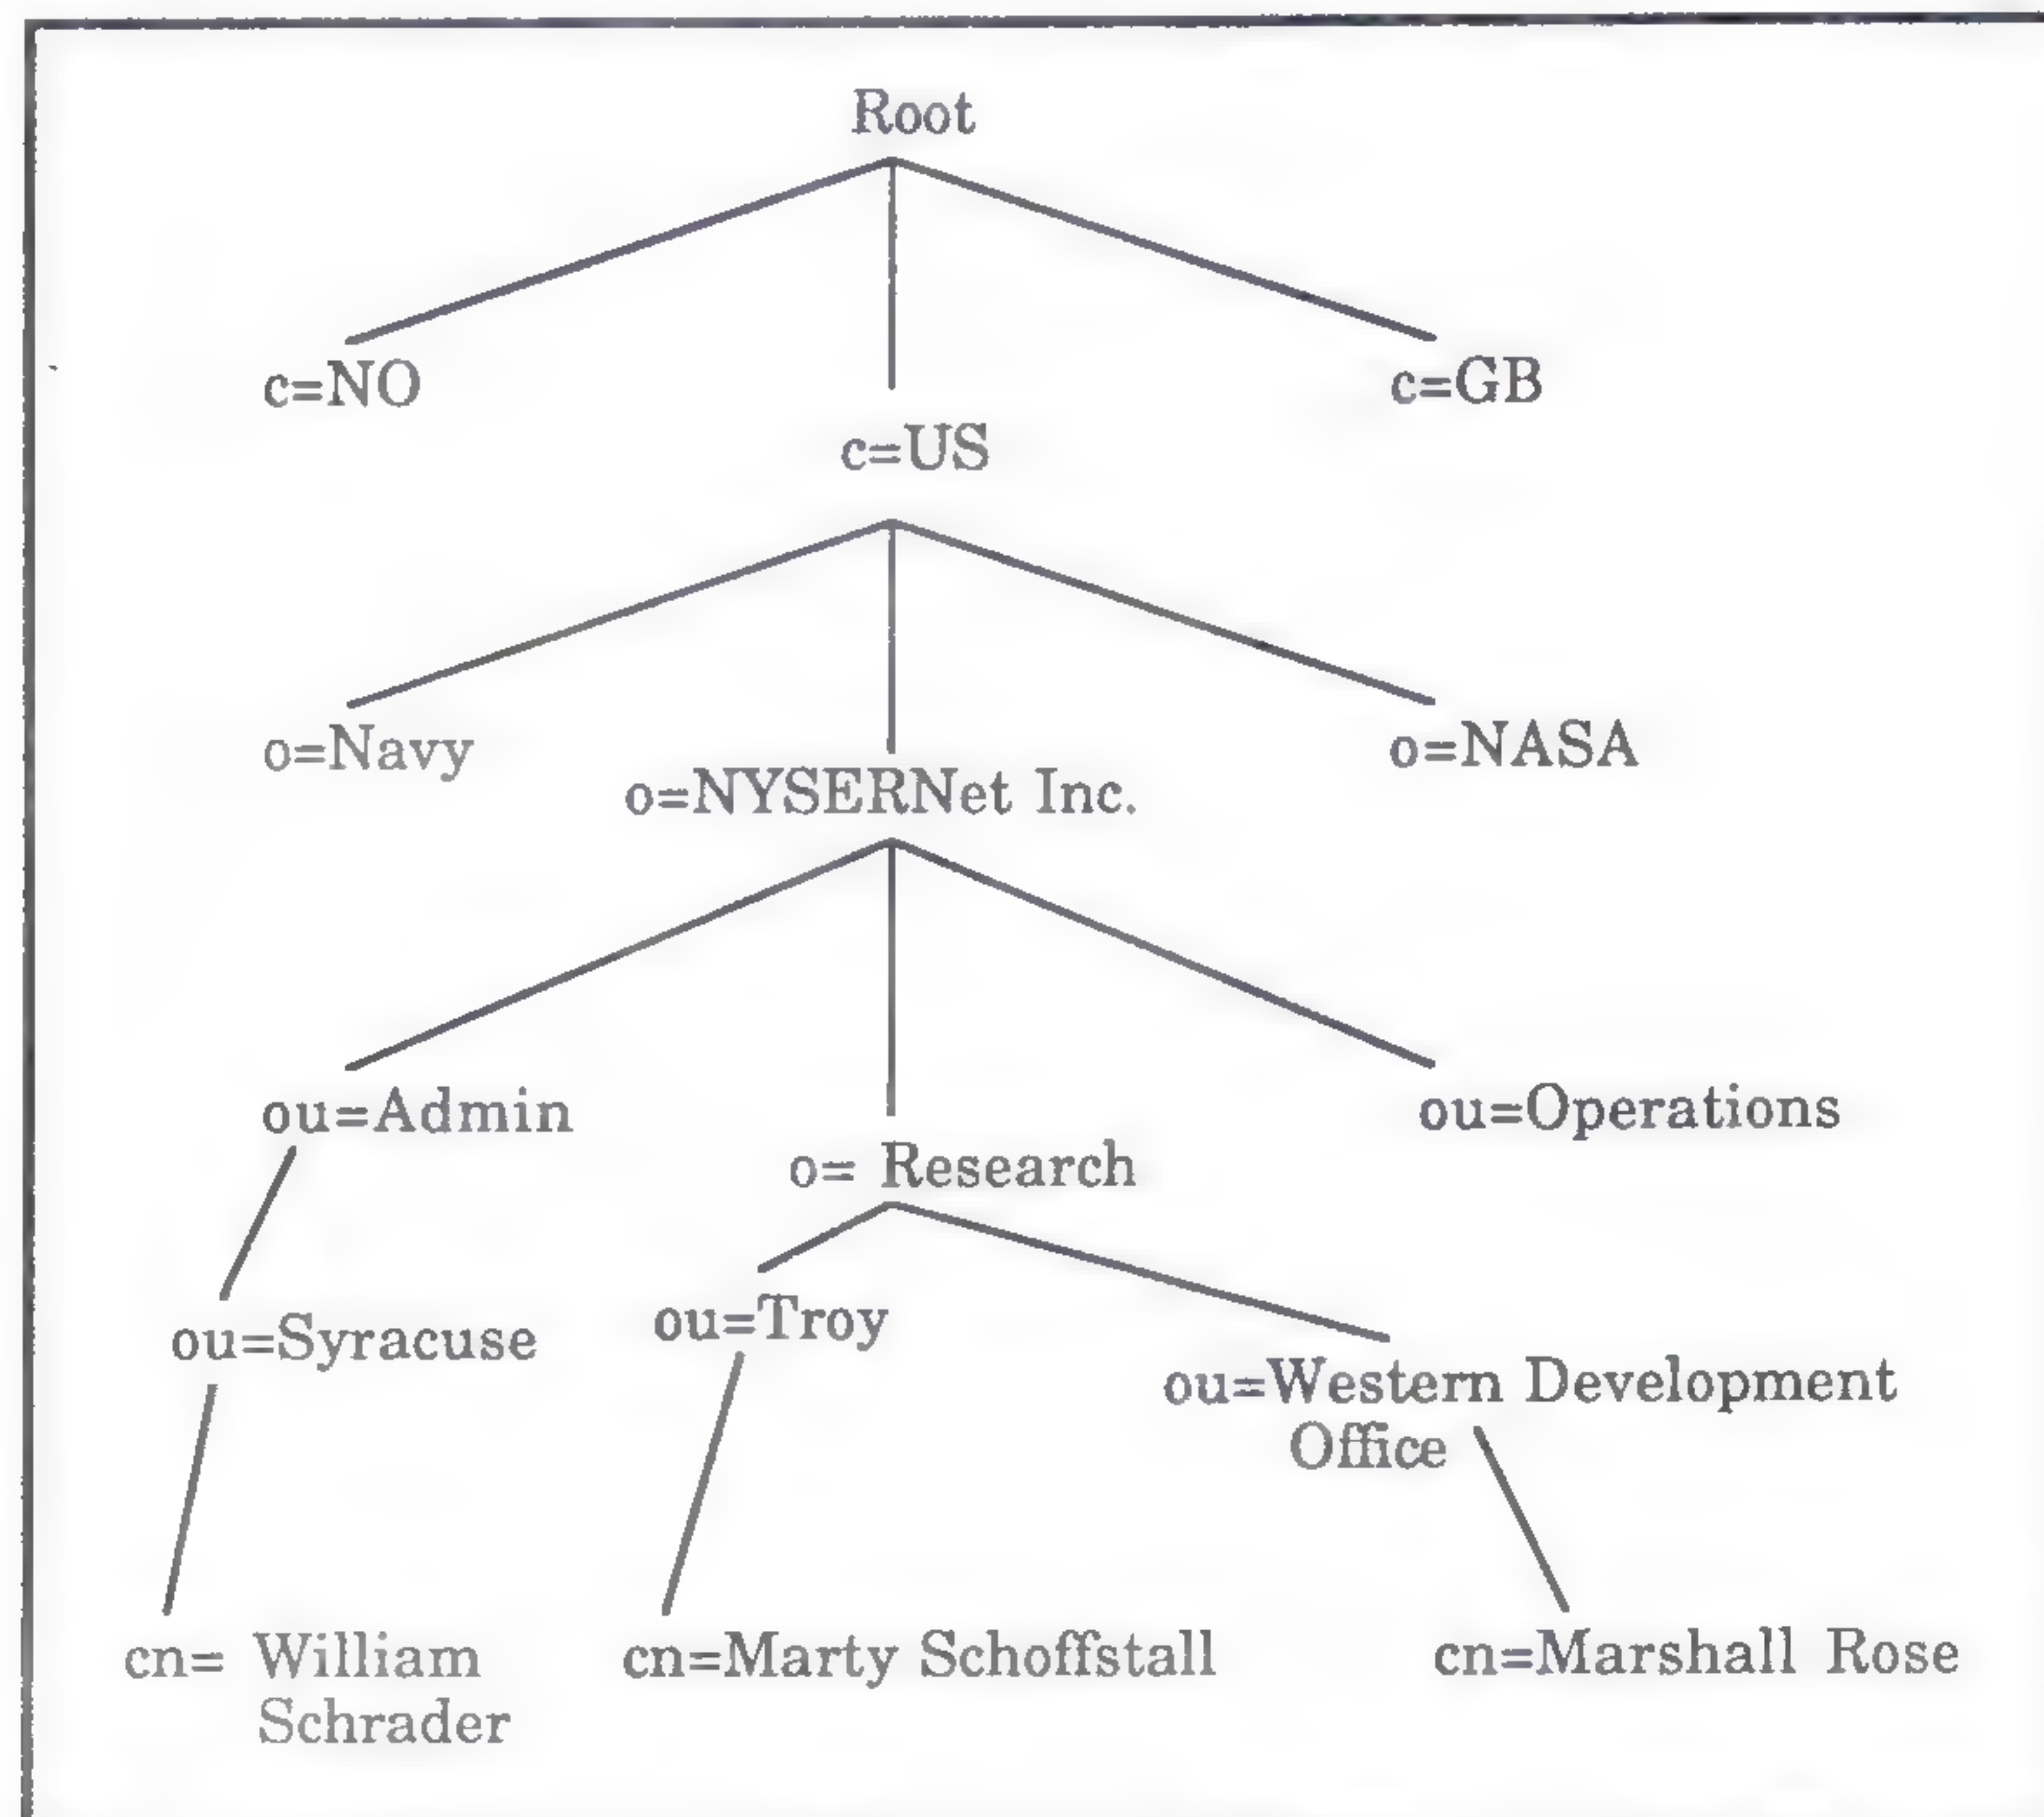
Finally, most sites maintain their own local white pages service which do not interoperate with the existing facility. This leads to at least two sets of user interfaces, procedures, programs, and so on.

It should be clear that any replacement facility must not only provide (at least) equivalent functionality to WHOIS, but must also address all three of these deficiencies. This replacement should be based on a standard distributed directory service model and the OSI Directory Service is the best available candidate.

The OSI Directory

The OSI Directory is designed to provide for the management of names and associated attributes. It is intended to provide a wide range of services, including both white pages (name-based searching) and yellow pages (attribute-based searching). [2]

The OSI Directory is structured in the form of a hierarchical tree. Each object in the tree has a *distinguished name*, which uniquely identifies it. Associated with each object is one or more *attributes* and possibly one or more child objects. An attribute might be something like a surname, a telephone number, an electronic mail address, a job title, or a textual description. Each attribute consists of a name and one or more values. Based on the name of the attribute, the value(s) are strongly-typed. This structure allows unambiguous understanding of the attribute, regardless of the program accessing the Directory. The OSI Directory standard defines several kinds of attributes along with their associated data types. In addition, users of the Directory may define additional attributes of their own.



Distinguished Names in the White Pages

White Pages service

In writing the software for the pilot, NYSERNet kept the user interface as close as possible to the original WHOIS interface. This is desirable since a large portion of the community has used WHOIS and is familiar with this style of interaction. Nevertheless, using the OSI Directory to provide the white pages, rather than a centralized database, led to three ramifications on the white pages service.

continued on next page

NYSErNet Sponsors White Pages Pilot (*continued*)

First, entries in the new service are uniquely identified by their distinguished name. Under the WHOIS service, a short string, such as "MTR" could be used. For larger communities, it is necessary to structure the names so as to delegate naming authority and de-centralize management. Thus, a unique handle in the white pages service is something like:

```
c=US@o=NYSErNet Inc.@ou=Research@cn=Marshall Rose
```

which is read left-to-right as:

```
country is US
organization is NYSErNet Inc.
organizational unit is Research
common name is Marshall Rose
```

While this is longer than a three or four letter acronym, it's the price one pays for a *global* naming scheme.

Second, searching is over portions of the white pages data. For a centralized service, it makes sense to exhaustively examine all entries when searching. For a distributed facility, this is too resource intensive. So, an interactive process is used, similar to the telephone white pages.

There are several telephone books, one for each particular geographical area. As such, before you can find someone's entry, you must first find the area in which they are listed. Because information on organizations is kept in the Directory, users first query the white pages to find the correct area; they then focus the white pages on that area. The pilot project software contains special heuristics which make this searching very efficient. In most cases, the correct area is found quickly which permits an intensive search to be focused accordingly.

Third, information in the white pages is highly structured. This allows several different programs to store and retrieve information and maintain consistent understanding of the value. For example, in addition to searching on a name or mailbox, using the Directory, one could just as easily search by telephone number. This structuring also allows different white pages interfaces to be built. NYSErNet currently supplies a text-based interface, and has plans to develop an X Windows interface.


An example

A simple example illustrates these concepts. Suppose a user is interested in finding information about someone called "Goodfellow" at someplace called "Anterior." The user invokes a user-interface to the White Pages called "fred," and issues the query:

```
whois goodfellow -org anterior
```

which means: "Tell me about anyone named 'goodfellow' at anyplace called 'anterior.'" The user-interface first determines which organizations have "anterior" in their name. Depending on user-options, this may use simple wild-carding, or a complex Soundex (phonetic) algorithm for matching. In this example, only one organization is found.

The user-interface then focuses its search on this organization asking about someone named "goodfellow." Again, depending on user-options, various matching algorithms might be used, on either the person's full name or surname only. As shown here, only one person is found:



```
% fred
fred> whois goodfellow -org anterior
Trying @c=US@o=Anterior Technology ...
Geoffrey Goodfellow (2) Geoff@Fernwood.MPK.CA.US
    aka: Geoffrey S. Goodfellow
```

```
President
    POB 1206
    Menlo Park
    California 94026-1206
```

```
Telephone: +1 415 328 5615
FAX:        +1 415 328 5649
TELEX:      number: 650 103 7391, country: US, answerback: MCI UW
```

```
Mailbox information:
MCI-Mail: Geoff
Internet: Geoff@Fernwood.MPK.CA.US
UUCP: fernwood!Geoff
```

```
Drinks:    chilled water (and lots of it)
```

```
Handle:
@c=US@o=Anterior Technology@ou=Corporate@cn=Geoffrey Goodfellow (2)
Modified: Fri Jul 21 11:41:27 1989
fred>
```

The first line shows the name of the person ("Geoffrey Goodfellow") along with an abbreviation for the person's handle ("2"), and the user's Internet mailbox. After this, any other names for the person is shown, in this case there is one, which includes the person's middle-initial. After this, the usual postal, telecommunications, and e-mail addressing information is shown. Next, the person's favorite drink is displayed ("chilled water"), along with the user's full handle. Finally, the date and time of the last modification to this entry is shown, to give the user an idea of how current the information is.

It should be noted that only some of this information is textual. In many cases, the attributes are complex binary structures, such as the facsimile image shown above. The user-interface is responsible for displaying this information in a pleasing fashion.

The NYSERNet Pilot

Participation in the NYSERNet White Pages Pilot Project is strictly voluntary—the pilot project is a "grass roots" effort, both to understand the white pages service desired by users along with the limitations of the OSI Directory in providing those services.

The primary goal of the pilot project is to encourage organizations to use the OSI Directory to store infrastructural information about their personnel. (Note that this does not require the elimination of existing mechanisms, such as internal telephone directories.) In addition, organizations are encouraged to maintain their own portion of the Directory tree. (For resource constrained sites, NYSERNet will offer maintenance service, just as it has been doing for the Domain Name System for the past three years.)

continued on next page

NYSErNet Sponsors White Pages Pilot (*continued*)

Another goal of the pilot project is to use the same programs and tools to access both global and local white pages information. As a part of this, new applications which might make use of the white pages service, such as private mail, will be encouraged.

When a NYSErNet member joins the pilot, it completes two phases. The first phase focuses on collecting data for the pilot project. For some participants, this will be completed in less than a week; for others, a month or so may be required. Since each organization usually already maintains this information in some local format, this task consists primarily of preparing the data for inclusion in the OSI Directory.

The second phase focuses on offering the service to the pilot user community. NYSErNet is providing the software to run the OSI Directory to all sites in source form, along with user and administrator guides. It is anticipated that all participants will have entered the second phase prior to the INTEROP™ 89 conference and exhibition, October 2–6. The NYSErNet booth at INTEROP will provide access to the pilot for demonstration purposes.

The pilot project will run until June 1, 1990. At this time, the pilot project will be evaluated. If successful, the membership to the pilot project may be expanded beyond NYSErNet, with Phase I being reactivated on a larger scale. Most likely this will also result in other applications, such as a window-based user interface being fielded.

Initial Membership

By the end of July 1989, the following organizations had received permission to participate, or were participating in the pilot project:

Alfred University	Polytechnic University
Anterior Technology	Rensselaer Polytechnic Inst.
Clarkson University	Rockefeller University
Columbia University	Syracuse University
Eastman Kodak Company	SUNY Albany
Engineering Information, Inc.	SUNY Buffalo
NASA	SUNY StonyBrook
Navy	University of Michigan
New Mexico State University	University of Pittsburgh
New York University	University of Rochester
NYSErNet, Inc.	

Not all of these are NYSErNet sites. Indeed, several organizations outside of NYSErNet have expressed an interest in joining the pilot project. Such petitions are decided on a case-by-case basis.

Conclusions

A white pages service has the potential to unify the management of the infrastructural information that is vital to networking. By sponsoring a pilot project, in addition to offering a valuable service to the user community, NYSErNet hopes to gain vital administrative and operational experience.

References

- [1] ISO/CCITT, "Information Processing Systems—Open Systems Interconnection—The Directory," ISO 9594-1-8, CCITT X.500–X.521, 1988.
- [2] Benford, S. "Components of OSI: The X.500 Directory Service," *ConneXions* Volume 3, No. 6, June 1989.

A Letter to the Editor

Dear Ole,

I found Ms. Nancy E. Hall's article in your July issue of *ConneXions* quite helpful. ["Components of OSI: The Transport Layer," Volume 3, No. 7]. It's too bad she had not written it before the National Research Council comparison of TP4 and TCP—it would have been very useful input.

Ms. Hall asserts that the UDP does not offer a *Quality of Service* parameter. While this is, strictly speaking, correct, the lower level IP layer does have several such parameters which are accessible, in theory, to the UDP level. It is arguable whether much is done with these parameters, so her assertion may be, in practice, correct, if technically inaccurate.

On page 7, Ms. Hall mentions classes 0 through 4 of the OSI TP protocol. In fact, there is a fifth class being defined. It is my understanding that this new class is targeted at the same general performance capability as Dr. Greg Chesson's *eXpress Transport Protocol* (XTP). I do not know the current status of Class 5 TP, but believe it is unlikely to have progressed beyond the draft proposal stage.

Finally, on pages 7 and 8, Ms. Hall lists classes 0 through 4 of TP and gives a terse description of each. In practice, I believe we are finding mostly implementation and use of TP0, TP2 and TP4. The last one more in the U.S. where there is more concern with datagram LANs. Indeed, The Europeans seem to be concentrating on either Class 0 (above X.25) or Class 2, while the U.S. seems more interested in connectionless modes supporting Class 4. Interworking among Classes 0, 2 and 4 of TP is not a trivial matter, as Ms. Hall points out at the bottom of page 8.

You deserve great credit, Mr. Editor, for your hard work. *ConneXions* is a useful publication and I have often turned to it as a ready source of current information and, not infrequently, as a source of historically useful data (e.g., Robert Braden's Internet Timeline).

Cordially,

Vint Cerf
Vice President
Corporation for National Research Initiatives

We appreciate the clarification by Dr. Cerf. As the OSI protocols mature and become widely used, we will undoubtedly have more articles on this subject. As always, we welcome comments, suggestions and questions from all our readers. Send your letters to:

ConneXions
480 San Antonio Road, Suite 100
Mountain View, CA 94040

The Effect of the Internet Worm on Network and Computer Security

by Fred Ostapik, SRI International

Introduction

The Chinese have a saying which aptly describes our present age. Roughly translated, it is "May you live in interesting times." This saying is wonderfully ambiguous. But we do live in interesting yet ambiguous times. While this Information Age has produced major benefits through rapid, technological growth, it seems to have also spawned a dark side in which this same technology is used for ignoble purposes.

As technological tools evolve and grow more complex, they place greater power into the hands of the users. If these tools are designed and built well, and if the users understand how to use them, they can be of enormous benefit to society. On the other hand, misuse of technology and its tools, either by accident, or design can create increasingly more unmanageable chaos.

The Internet Worm

A striking example of such misuse is the case of the "Internet Worm." By exploiting deficiencies in the Berkeley UNIX electronic mail system and in other network utilities used by the Internet community, one person was able to disrupt the normal processing of thousands of host computers. Fortunately, the disruption was detected within hours of its occurrence and was eventually neutralized—an effort that required the cooperation of countless experts working throughout the United States who, ironically, used the very same network links that supported the transmission of the Worm.

The media focus on worms and viruses may present the erroneous impression that network and computer system security mainly addresses problems arising from external intrusions. Another factor contributing to this misconception is the confusion in people's minds regarding the term "security." The usual definition is linked with the government's idea of security as protection from compromise external to an organization. In reality, this limited concept only creates a false sense of security. There may be other, possibly far more devastating, villains. Consider the recent trend of spies originating from the very governmental agencies charged with protecting us from such miscreants.

A more meaningful definition of network and computer security involves the concept of protecting data, systems, and networks from any compromise at all. After all, the net result is the same if someone's twenty years of research data is lost, whether accidentally or by deliberate design.

Threats

Many studies have been done to determine the vulnerabilities of information systems, and they tend to corroborate the conclusions that:

- Most vulnerabilities or compromises (over 80%) arise from internal sources. The vast majority of problems occur because of errors caused by honest employees who may be badly supervised, poorly trained, or under a great deal of stress.
- Another source of security risk is employees who are dishonest. The main motivation for their actions is greed or some other impetus toward personal gain.

- Disgruntled employees form the next highest group. They are generally seeking some form of revenge.
- Natural catastrophes and disasters cause 5% to 15% of the problems.
- All other causes make up less than 5% of the incidents. This group includes compromises caused by external threats.

Even though external threats make up a very minor percentage of all disruptions, the notoriety given by the media to the more flamboyant intrusions may provide the incentive for others to make similar attempts. Furthermore, because of the relaxed Internet philosophy of encouraging the open exchange of ideas, the Internet continues to be very vulnerable to unscrupulous hackers.

The "Internet Worm" incident may have produced some indirect benefits, however. The great attention drawn to the problems of maintaining network and information systems security may result in a concentrated effort by government, research organizations, and commercial groups to address some of the important issues. Certainly the "Worm" has created full employment among security consultants.

GAO Report

The *General Accounting Office* (GAO) recently issued its own "Worm" report. It concluded that the primary problem was the absence of a government "security focal point" that could spearhead the counterattack against catastrophic intrusions. It recommended that an interagency group, to be appointed by President Bush's science advisor, serve as this "focal point."

[Ed.: The GAO report is available by anonymous FTP from the NSF Information Services host `nis.nsf.net` (35.1.1.48) in directory `nsfnet`. Log in as user "anonymous," with password "guest." Change directory to `nsfnet` ("cd nsfnet"), and get the file `GAO_RPT.TXT`].

Internet experts, however, have been quick to point out that an "Internet Worm" antidote was produced within hours of initial intrusion, without benefit of a government focal point. They worry that the involvement of another agency would only add more bureaucracy and delay. However, they do feel that such an agency could serve as a resource for long-term solutions to some of the Internet security problems.

Security Coordination Center

An informal *Security Coordination Center* (SCC) has been created at SRI International to be a clearinghouse of information on security problems affecting primarily the hosts on the MILNET and ARPANET. (It is in its final stages of being formally established by the Defense Communications Agency.) During the "Internet Worm" incident, SRI worked with the key players at Berkeley and other sites to gather the information necessary to fight the problem, to have this information validated for accuracy, and to distribute this information to ARPANET and MILNET users. Since then, a number of other potential security problems have been uncovered which affect sites on the Internet, and SRI has been actively involved in keeping these sites informed of the resolution of these problems.

continued on next page

Network and Computer Security (*continued*)

SRI works closely with other organizations, such as the *Computer Emergency Response Team* (CERT), to insure that the information it distributes has been validated to the utmost possible extent before it is disseminated. SRI is also conducting research into various aspects of computer and network security. It has developed prototype intrusion detection and auditing processes that attempt to address the problems involving compromises originating both inside and outside an organization.

Kerberos A group at MIT has successfully implemented *Kerberos*, an authentication system designed to work in an open network computing environment. It provides a mechanism, a trusted third-party authentication service, to verify users' identities when they wish to use network services such as electronic mail.

Vendors Vendors are also becoming increasingly attuned to the concerns about the robustness of their products. They seem to be giving greater priority to correcting security problems in their installed systems when these problems are reported by their customers. Vendors are hearing the demands of the marketplace for the development of products which enhance the protection of the customer's information systems. For example, RSA Data Security Inc. has devised a system that will permit users to send encrypted messages over a network in "digital envelopes" that can only be accessed by the addressee. The contents will contain a "digital signature" that cannot be forged. This technology is targeted for implementation on the Internet.

Conclusion In conclusion, the "Internet Worm" incident may have been the watershed event that has produced a decisive change of attitudes, a toughening of attitudes among the public towards computer penetrations. The definition of computer crime is being extended through legislation and the judicial process. No longer is an unauthorized romp through someone's data files considered an innocent pastime for clever young people who wish to be hired as consultants by the very institutions they have violated. Without a doubt, the seriousness of the "Internet Worm" incident has been injected into the public psyche. But whether this injection will have the effect of an inoculation—of building mechanisms for immunity against future worm and virus attacks—remains to be seen.

The issues discussed in this article will be further explored in the upcoming *InfoSec* seminar, see the following page. —————>

[Ed.: See also "Long-Term Implications of the Internet Worm" in *ConneXions*, Volume 3, No. 4, April 1989.]

FRED M. OSTAPIK is a Senior Research Analyst at the Network Information Systems Center at SRI International. His tasks include the design and implementation of audit trail systems for the Defense Data Network (DDN). He led the effort to produce NAURS, a world-wide network audit trail and usage data collection and reporting system, used to produce billing, trend analysis, and capacity planning reports for the ARPANET and MILNET. Previously he was the Director of Computing Services and Institutional Research at San Francisco State University, and Manager of Computing Services at the University of Wisconsin-Milwaukee. Fred holds an M.S. in Computer Science, and a B.S. in Applied Mathematics and Physics from the University of Wisconsin-Madison.

Upcoming Events

InfoSec™ 89—Practical Perspectives on Computer and Network Security will be held November 28-30, 1989 at the Desert Princess Doubletree Hotel in Palm Springs, California

The seminar is sponsored by Advanced Computing Environments and SRI International.

InfoSec is a unique, three day seminar dedicated to improving awareness of network and computer security issues among MIS Managers and Networking Professionals who are not experts in the area of security. InfoSec provides a one-day tutorial and two days of technical seminars presented by leading experts from the commercial, government, vendor and research communities.

InfoSec will bridge the gap between theory and practice and will provide you with detailed and up-to-date knowledge on how to maintain and increase security at your computer and network system installations. Leading vendors will host peer-to-peer discussions featuring their technical specialists in computer and network security.

Tutorial The first day provides a comprehensive tutorial, led by Steve Walker of Trusted Information Systems. The tutorial covers security concepts, goals and objectives, currently available technologies, and introduces the attendee to the security problems which must be solved now and in the future.

Sessions On Day 2 and Day 3 we present six Technical Seminars which analyze the state of the art in computer and network security. The sessions will be chaired by leading experts from the research, government, vendor and user communities. The focus will be on the practical options available to maintain and increase computer and network security. Session titles include:

- Technology Assessment and Future Prospects
- User Perspectives
- The Role of Government
- Vendor Perspectives
- Legal, Social and Political Implications of Computer and Network Security

A detailed InfoSec Advance Program will be mailed to you in the near future.

Don't miss this opportunity to learn more about the current state-of-the-art in computer and network security.

For more information on InfoSec, or to register, call Advanced Computing Environments at 415-941-3399 or FAX at 415-949-1779.



AppleTalk versus IP

by Greg Minshall, Kinetics Inc.

Introduction

This article discusses some of the differences between “plug-and-play” networks like *AppleTalk*, and IP networks which require a great deal of configuration and expertise on the part of the network operator.

Setting up an AppleTalk network

Simple networks need no installation. These networks consist of exactly one cable (the *LocalTalk* wire) connecting a small number of end systems (around 20). The end systems dynamically acquire their own node number (something equivalent to the low order byte of their IP address).

More complicated networks require one or more AppleTalk routers. Configuring these requires inventing one or more zone names (one is fine, no matter how many networks are interconnected; more is a performance/convenience issue), and inventing some unique network numbers (16 bit integers). (There exists routers which allow the user to bypass even these steps by offering an “auto configuration option”). No change is needed to the non-router nodes.

All routers use the same routing and resource location protocols, so there is “no” user configuration in this area.

If the network number needs to change, the administrator needs to reconfigure the AppleTalk routers connected to the affected network and then restart them.

Setting up an IP network

In contrast, setting up an IP network is more complicated: Each machine must be assigned its own IP address (either set on the machine itself or in a file like */etc/bootptab*). These addresses must obey a rather complicated set of rules which are not easily explained to the casual listener. (“Complicated” in the same sense that last year’s 1040EZ tax form was seen as “too complicated”). Even when setting up the most simple network (one wire, just a few machines), these addresses must be assigned by hand.

When adding a router to a previously isolated network, it may be necessary to change each host (add a static route or start up */etc/routed* for instance) before they will become aware of the new router. Each router must be configured to speak the correct routing protocol(s) in use by end systems and by other routers with which it needs to interoperate.

At some point in the evolution of a network there will arise the need to split the previous network into *subnets*, causing endless debate and hair-pulling by end-users as to “what does it all mean, after all?” (And when you come down to it, effectively, no one actually knows!)

If the network number needs to change for a given network, *each* end system’s configuration needs to change, and each end system restarted (in one way or another) for the change to take affect. Woe to the network on which some end systems know the new number and others the old.

Scale

Now, let’s be fair. It is very likely that AppleTalk will not scale up the way IP has already scaled up. There are a lot of AppleTalk machines out there, and there are a lot of AppleTalk networks.

But, I would bet that the largest AppleTalk internetwork has less than 200 AppleTalk networks on it. Additionally, I doubt that the geographical area of very many AppleTalk networks is more than twenty five miles or so.

In contrast, the TCP/IP Internet spans the globe and has (I suppose) thousands of networks attached to it. Part of the initial thrust behind TCP/IP was "wide area" (and part was "heterogeneous networks").

However, how many of "us" actually benefit from the wide area nature of IP? Certainly those that use the TCP-IP mailing list, FTP files from around the country, etc. Think, though, of all the people at Motorola in Tempe, Arizona (is there such a site?) who are using `rlogin`, `ftp`, `rsh`, etc., *right this minute*, without ever accessing a computer farther from their desks than three miles. Do they *really* benefit from the "wide area" nature of IP? They certainly benefit from the heterogeneity; but do they really benefit from the "wide area" nature? I suspect not.

**The vendor's
viewpoint**

I come to this question from a commercial point of view. We sell (in some sense) AppleTalk networks. We also sell IP networks (again, if you think of things in a certain way). Our users are (more or less) able to set up AppleTalk networks with no intervention from us. No intervention means no phone time with our tech support personnel (not to mention no phone time with our engineering staff). This means that supporting AppleTalk networks is inexpensive for us.

However, our users are not able to set up IP networks without a fair (to substantial) amount of intervention with our technical support staff (and, sometimes, some phone time with our engineering staff).

Think of it from our point of view. We have a certain reputation within our customer base. We value that reputation; it is part of our business assets, if you will. Those customers of ours who are setting up AppleTalk networks allow us to provide them a reasonably good product at the time of sale, and then cost us very little in post-sales support to keep them "happy," i.e., to maintain our reputation as providing good products and good support. However, to keep our IP customers happy, we need to expend a larger amount of money. And, even in doing this, we lose because the "casual listener" really believes that things must be simpler, and that therefore there *must* be something wrong with the product we are selling or with what we are telling them to do.

Additionally, think of it from the point of view of a company providing application level products which make use of a network. ("Application level" disallows things like IP routers, TCP for VMS, etc.) Manufacturers can easily sell such products in an AppleTalk environment. This is because they do not need to instruct the end-user on how to set up an AppleTalk network. They just say:

"Install the server wherever, give it a *name* and run the client and (if asked) select the zone in which the server was installed from the supplied (and automatically determined) list of zones. Then select the specific server by name from the supplied (and automatically determined) list of our servers running in the selected zone."

continued on next page

AppleTalk versus IP (*continued*)

In fact, part of the power of AppleTalk networking is that the user who installs the client doesn't need to know *exactly where* the server is located. The protocol stack allows the client to acquire a list of "named collections of networks" (known as *zones*), and then look in any or all of these zones to find available servers of any specific type. This is powerful because it frees the end-user from knowing where the server was installed; it also frees *manufacturers* from describing network numbers, node numbers, routing protocols, etc., to the end user. This latter reason is of great economic benefit to the manufacturer.

The problem with IP

For a program running in an IP environment, the manufacturer needs to first of all define a bunch of terminology (subnet mask, IP address, RIP, etc.), then lead the user through the task of defining where the server is installed, then configuring the client with the address of the server, explaining what the customer should do if things don't seem to be working out ("`netstat -r`," etc.) And, again, this is the case *even if* this entire customer installation is contained within two floors of one building. (None of this is to say that there *aren't* occasionally problems with AppleTalk networks that take experts some time to track down. We're just talking about orders of magnitude here).

In some sense, this is one of the main problems with IP as we currently define it. It requires *too much knowledge from too many people*.

Above, I mentioned the cost to us, a networking company, of supporting IP networks. We are, however, a networking company. We see a clear need to support IP networks (and, clearly, many of our roots are in IP networks; and IP networks, even with the resulting customer support burden, provide a fair portion of our income). However, from the point of view of a company manufacturing application level products, the cost to support IP networks must seem unsupportable.

The fact is that there are more application level products which rely on networking sold in the AppleTalk arena than there are in the IP world if we leave out products sold by networking companies (i.e., companies which are selling TCP/IP protocol support).

Why should this be of concern to the TCP/IP community? Is this not just of concern to rabid merchants like Mr. Minshall? I think that most of us would like to see IP used by as large a group as possible. There are many reasons for this:

- If even our doctor's office is wired for IP, our own knowledge of, and involvement in, IP technology is going to provide us with a larger ego boost than we currently receive.
- Our earning power (either in 9-5 jobs, in consultancies, or in research grants) is, most likely, directly tied to the number of IP installations.
- We get much ego satisfaction in seeing programs we wrote (either in a public domain context, such as 4.3BSD, or as products, such as the K-STAR software of a Kinetics FastPath) distributed and used widely.

Additionally, trying to ease the installation of IP networks and IP applications is an interesting research problem (but maybe not the most interesting in the world).

Making IP networks simpler to install

Above I mentioned that two of the design goals behind the IP “experiment” were “wide area networking” and “heterogeneity.” These two goals have been met admirably. I would suggest, however, that we might think about making “ease of installation” (of both networks and network-using applications) a goal. In fact, it would be my opinion that “ease of installation” should take precedence over “wide area networking.”

Now, hang on a second. I am *not* suggesting that we shouldn’t be able to FTP half-way around the world. However, I would be willing to make the task of setting up (and maybe using) this ability a bit harder if that was necessary (and it probably wouldn’t be) to meet the “ease of installation” goal. The fact is, however, I believe that we could (if we set our collective minds to it) meet the “ease of installation” goal, and still have a network which scales up (though functionality may decrease the higher “up the scale” you go).

There has been ongoing work in the TCP/IP community trying to make things easier for end users. At least two computer companies have attempted to provide TCP/IP products which can be configured to dynamically acquire an IP address at boot time (ideally, this capability would not need to be configured, but would, rather, happen by default). Additionally, the existence of protocols such as BOOTP (a protocol to allow TCP/IP end systems to retrieve their configuration information from a centralized database) allow for the design of schemes for dynamic address assignment by a BOOTP server. Thus, the problem of acquiring an IP address is probably close to resolution. [Ed.: See *ConneXions*, Vol. 2, No.10, October 1988, p. 14].

The second problem, of registering, listing, locating, and accessing network resources, with little or no prior knowledge, is a bit less tractable at the current time. However, research groups, as well as at least one computer company, are working at trying to solve this problem.

What’s wrong with AppleTalk?

After taking what might seem like such a large number of unwarranted pokes at IP, it seems only fair that I point out some of the problems with the AppleTalk protocol suite:

- There are only 65,536 AppleTalk networks available, as opposed to IP’s $128+32,768+2^{23}$. The AppleTalk address space is 2^{24} , as opposed to IP’s 2^{32} . Thus, there is a resource shortage.
- The scheme for naming zones is not hierarchical, and a zone name is limited to 32 characters. This imposes a flat naming space which almost certainly would not scale to Very Large Networks (VLNs).
- The checksum procedure used by AppleTalk’s DDP protocol (the equivalent of IP *and* UDP) does not allow for checksumming of the (routing portion of the) header at inter-network routers (while IP’s does).

continued on next page

AppleTalk versus IP (*continued*)

- DDP allows only a maximum of 15 hops before a packet is discarded.
- There is no provision within the AppleTalk protocol for hierarchical routing (something IP has accomplished, within some view of things, via subnets). Thus, each internetwork router needs to know the entire network topology. Again, this does not scale up well to VLN's. (To be truthful, though, the subnet scheme doesn't really solve this problem for the IP community; it has just alleviated the problem for the time being).
- There is no way of reserving a specific network entity name (the entity looked up by client applications when trying to find servers) to ensure that a given physical server is able to keep the same name each time it is restarted (the names need to be unique, and there is nothing to keep an undergraduate in a dorm from naming the local printer "Chancellor's Office" if the printer in the chancellor's office is powered off at the time).
- There is no equivalent of an "ICMP redirect," thus causing many packets (originating on networks with multiple routers) to take an extra hop before reaching their destination.

Clearly, the problems with the AppleTalk protocol stem from its limited vision of how the large an AppleTalk internetwork is going to be—problems which IP has (for the most part) solved (or at least deflected). (This limited vision was, more than likely, one of the design tradeoffs made in the design of the AppleTalk protocols—not the limited vision of the original designers).

GREG MINSHALL is a programmer with Kinetics—a division of Excelan Inc., a subsidiary of Novell, Inc.—specializing in providing networking products for Apple Macintosh computers. Greg also spends a small portion of his work time at UC Berkeley.

New RFC on Network Management

RFC 1109: Report of the Second Ad Hoc Network Management Review Group is now available from the Network Information Center in the online library at NIC.DDN.MIL.

This RFC reports an official Internet Activities Board (IAB) policy position on the treatment of Network Management in the Internet. This RFC presents the results and recommendations of the second Ad Hoc Network Management Review on June 12, 1989. The results of the first such meeting were reported in RFC 1052. This report was approved and its recommendations adopted by the IAB as assembled on July 11–13, 1989. Distribution of this memo is unlimited.

RFCs can be obtained via anonymous FTP from NIC.DDN.MIL, with the pathname RFC:RFCnnnn.TXT (where "nnnn" refers to the number of the RFC). The NIC also provides an automatic mail service for those sites which cannot use FTP. Address the request to SERVICE@NIC.DDN.MIL and in the subject field of the message indicate the RFC number, as in "Subject: RFC 1109." Hardcopy RFCs may be ordered by calling the NIC at 1-800-235-3155 or 1-415-859-3695.

The Internet Activities Board and its Task Forces

Introduction

The Internet, begun as a DARPA-funded experiment in inter-networking, has become the prototype national research network, forming a major fraction of the networking infrastructure for the research community and, increasingly, the government and business sectors. The TCP/IP technology has been propagated even more widely into business and industry.

Networking and internetting are now entering a new era with the emergence of extremely high speed packet switching technology and with the steadily growing availability of OSI software. The decade of the 1990's is likely to prove as revolutionary in network technology development as the ARPANET was almost twenty years ago.

The *Internet Activities Board* (IAB) functions as a Board of Directors for the Internet. The IAB itself set technical policy and standards for the Internet protocols and architecture, and reviews the work of its two major task forces: the *Internet Engineering Task Force* (IETF) and the *Internet Research Task Force* (IRTF).

Representatives of the government agencies that sponsor important segments of Internet research and infrastructure are members of the *Federal Research Internet Coordinating Committee* (FRICC). The FRICC is concerned with the present Internet and the future National Research Network. The IAB and the FRICC work closely together.

The IAB

The IAB is an independent committee whose members generally share a long-term involvement in, and responsibility for, Internet design, engineering, and management. The IAB members were chosen for the specific roles they play (e.g., the chairs of the IETF and IRTF), to represent major groups in the Internet community (e.g., national network, vendor, government, and international groups), and for specific areas of expertise (e.g., security). They are deeply committed to making the Internet function effectively and to evolving the Internet to meet a large scale, high speed future. All IAB members are required to have at least one other major role in the Internet community in addition to their IAB membership.

Members

The IAB chair serves for a period of two years. The current IAB chair is Vint Cerf of NRI, who has been called the "Father of the Internet." In 1974, he co-authored with Bob Kahn the original paper on the Internet architecture. As a program manager at DARPA he directed the research effort that resulted in development of the TCP/IP protocol suite, and directed the construction of the prototype of today's Internet.

The IAB membership is currently as follows:

Vint Cerf	Chairman
Dave Clark	IRTF Chairman
Phill Gross	IETF Chairman
Jon Postel	RFC Editor
Bob Braden	Executive Director
Hans-Werner Braun	NSFNET Liaison
Barry Leiner	CCIRN Liaison
Dan Lynch	Vendor Liaison
Steve Kent	Security

continued on next page

The IAB and its Task Forces (*continued*)

The IAB typically meets four times a year, to supplement its work by electronic mail. Whenever possible, and at least once a year, the IAB plans to hold its meetings in conjunction with an IETF meeting.

In more detail, the IAB performs the following functions:

- Sets Internet Standards
- Manages the RFC publication process
- Reviews the operation of the IETF and IRTF
- Performs strategic planning for the Internet, identifying long-range problems and opportunities
- Acts as external policy representative for the Internet community
- Resolves technical issues which cannot be treated within the IETF or IRTF frameworks

Internet Engineering Task Force

The Internet has grown to encompass a large number of widely geographically dispersed networks in academic and research communities. It now provides an infrastructure for a broad community of interest. Moreover, the family of Internet protocols and system components has moved from experimental to commercial development. To help coordinate the operation and management of the Internet, the IAB established the *Internet Engineering Task Force* with the charter to:

- Act as a clearinghouse to promote the exchange of information within the Internet community. This community includes Internet software and hardware developers, Internet operators and Internet research and development groups.
- Identify pressing and relevant short- to mid-range problem areas and convene Working Groups to develop solutions. Working Groups might deal with a wide range of Internet issues, such as operational management problems or protocol enhancements that would improve Internet performance or extend the range of application of the architecture.
- Report Working Group results and recommendation to the IAB and to the Internet community at large.

The Internet Engineering Task Force is a large open community of network designers, operators, vendors, and researchers concerned with the Internet and the Internet protocol suite. The work of the IETF is governed by a board known as the *Internet Engineering Steering Group*, or IESG. The chairman of the IETF and of the IESG is Phill Gross of NRI.

Working Groups

The work of the IETF is performed by subcommittees known as *Working Groups*. There are currently more than 20 of these. Working Groups tend to have a narrow focus and a lifetime bounded by completion of a specific task, although there are exceptions. The IETF is generally the major source of proposed protocol standards, for final approval by the IAB.

The IAB has delegated to the IETF general responsibility for making the Internet work, and, to the IESG in particular, responsibility for the resolution of all short- and mid-range protocol and architectural issues required to make the Internet function effectively.

The members of the IESG are, in addition to the IETF chairman, *Area Technical Directors* (ATDs)*. Each ATD has primary responsibility for one area of Internet engineering activity, and hence for a subset of the Working Groups. The ATDs perform a technical management function that is vital to the function and development of the Internet; they are selected not only for their technical expertise but also for their managerial skills and judgment.

For more information on the Internet Engineering Task Force, on attending meetings and proposing or joining Working Groups, contact Phill Gross, Corporation for National Research Initiatives. (gross@sccgate.scc.COM, pgross@NRI.Reston.VA.US, 703-620-8990).

Internet Research Task Force

The IAB is concerned with promoting research in networking, to develop the technology that the IETF will need; this is the job of the *Internet Research Task Force*, or IRTF.

In the area of network protocols, the distinction between research and engineering is not always clear, so there will sometimes be overlap between activities of the IETF and the IRTF. There is, in fact, considerable overlap in membership between the two. This overlap is regarded as vital for cross-fertilization and technology transfer. In general, the distinction between research and engineering is one of viewpoint and sometimes (but not always) time-frame. The IRTF is generally more concerned with understanding than with products or standard protocols, although specific experimental protocols may have to be developed in order to gain understanding.

The IRTF is a community of network researchers, generally with an Internet focus. The work of the IRTF is governed by a board known as the *Internet Research Steering Group* or IRSG. The chair of the IRTF and of the IRSG is Dave Clark of the MIT Laboratory for Computer Science. Clark played a seminal role in development of the Internet architecture, and for nearly 10 years guided the Internet development and deployment as chair of the IAB and as "The Internet Architect."

Research Groups

As befits their research goals, the IRTF and IRSG are organized in a much less formal manner than are their engineering counterparts. The IRTF is organized into a number of *Research Groups* (RGs), and the chairs of these RGs sit on the IRSG. These groups typically have 10 to 20 members, and each covers a broad area of research, pursuing specific topics, determined at least in part by the interests of the members as well as recommendations from the IAB and IETF.

* [Ed.: Most of the IETF Area Technical Directors are still to be appointed, we will bring you an update in a future issue of *ConneXions*].

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

CONNEXIONS

PUBLISHER Daniel C. Lynch

EDITOR Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President, National Research Initiatives.

Dr. David D. Clark, The Internet Architect, Massachusetts Institute of Technology.

Dr. David L. Mills, NSFnet Technical Advisor; Professor, University of Delaware.

Dr. Jonathan B. Postel, Assistant Internet Architect, Internet Activities Board; Division Director, University of Southern California Information Sciences Institute.

Subscribe to CONNEXIONS

U.S./Canada \$125. for 12 issues/year \$225. for 24 issues/two years \$300. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Charge my ☐ Visa ☐ MasterCard ☐ Am Ex Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

480 San Antonio Road Suite 100
Mountain View, CA 94040

415-941-3399 FAX: 415-949-1779

Back issues available upon request \$15./each
Volume discounts available upon request

CONNEXIONS